

Theoretical Mathematics & Applications, vol. x, no. x, 2013, x-x

ISSN: 1792-9687 (print), 1792-9709 (online)

Scienpress Ltd, 2013

A Consistency Result for Bayes Classifiers with Censored Response Data

Priyantha Wijayatunga¹ and Xavier de Luna²

Abstract

Naive Bayes classifiers have proven to be useful in many prediction problems with complete training data. Here we consider the situation where a naive Bayes classifier is trained with data where the response is right censored. Such prediction problems are for instance encountered in profiling systems used at National Employment Agencies. In this paper we propose the maximum collective conditional likelihood estimator for the prediction and show that it is strongly consistent under the usual identifiability condition.

Mathematics Subject Classification : xxxxx

Keywords: Bayesian networks, maximum collective conditional likelihood estimator, strong consistency

1 Introduction

Naive Bayes classifiers have proven to be useful in many prediction problems with complete training data. Here we consider the situation where a naive

¹Department of Statistics, Umeå School of Business and Economics, Umeå University, Umeå 90187, Sweden. e-mail: priyantha.wijayatunga@stat.umu.se

²Department of Statistics, Umeå School of Business and Economics, Umeå University, Umeå 90187, Sweden. e-mail: xavier.deluna@stat.umu.se

Bayes classifier is trained with data where the response is right censored. Such prediction problems are for instance encountered in profiling systems used at National Employment Agencies. A profiling system provides predictions of unemployment duration based on individual characteristics. To train such a system register data on individuals is used where unemployment durations as well as demographic and socio-economics information is recorded. Unemployment duration is then typically censored by the end of the observation period as well as exit from unemployment due to other reasons than employment, typically entrance into educational programs [1]. Naive Bayes classifiers have proven useful in many prediction problems with complete data [5], [3] and references therein. In this paper for this censored response case we propose the maximum collective conditional likelihood estimator and show that it is strongly consistent under the usual identifiability condition [4] whose notation is used in our proof.

Formally, consider a class variable X_0 –say unemployment duration– and a n –attribute random variable vector $X_{[n]} = (X_1, \dots, X_n)$ –individual features, where all the variables are discrete and finite. Note that continuous variables may be discretized making this framework more general. We assume thus that the state space of X_i is $\mathcal{X}_i = \{1, \dots, r_i\}$. We further assume that $(X_0, X_{[n]})$ forms a naive Bayesian network so that their joint density and conditional density of X_0 given $X_{[n]} = x_{[n]}$ are as follows:

$$p(x_0, x_1, \dots, x_n) = p(x_0) \prod_{i=1}^n p(x_i | x_0),$$

$$p(x_0 | x_{[n]}) = \frac{p(x_0) \prod_{i=1}^n p(x_i | x_0)}{\sum_{x'_0} p(x'_0) \prod_{i=1}^n p(x_i | x'_0)}.$$

Let the parameter space be $\Theta = \Delta_{r_0} \times \Delta_{r_1}^{r_0} \times \dots \times \Delta_{r_n}^{r_0}$ where $\Delta_t = \{(p_1, \dots, p_{t-1}) : 0 \leq p_i, \sum_{i=1}^{t-1} p_i \leq 1\}$ and for Δ_a^b , where b is the usual power. The interior of Θ is denoted by Θ° . In the following, we always assume that the true parameter is an element of Θ° . Note that, if this is not the case then the naive Bayesian network is degenerated in the sense that some variables (if binary) may vanishes as their state space will be reduced to singletons or some may appear with reduced state spaces. Then, the parameter of the naive Bayesian model is $\theta = (\theta_0, \theta_1, \dots, \theta_n) \in \Theta$, where $\theta_0 = (\theta_{x_0=1}, \dots, \theta_{x_0=r_0-1})$ and $\theta_i = (\theta_{i|x_0=1}, \dots, \theta_{i|x_0=r_0})$ such that $\theta_{i|x_0} = (\theta_{x_i=1|x_0}, \dots, \theta_{x_i=r_i-1|x_0})$ for $i = 1, \dots, n$. Since we are working on the non-Bayesian case we have

$$\begin{aligned}
p_\theta(x_0) &= \theta_{x_0} && \text{if } x_0 = 1, \dots, r_0 - 1, \\
p_\theta(r_0) &= 1 - \sum_{x_0=1}^{r_0-1} \theta_{x_0}, \\
p_\theta(x_i \mid x_0) &= \theta_{x_i \mid x_0} && \text{if } x_i = 1, \dots, r_i - 1, \\
p_\theta(r_i \mid x_0) &= 1 - \sum_{x_i=1}^{r_i-1} \theta_{x_i \mid x_0}.
\end{aligned}$$

Note that the above marginal and conditional densities are over-parameterized, i.e., when we write, for example, the density of X_0 , $p_\theta(x_0)$ the parameter θ also contains irrelevant components in addition to relevant ones to determine the probabilities of X_0 .

Suppose we have a random sample with $N = N_1 + N_2$ number of data cases on the random variable vector (X_0, X_1, \dots, X_n) , and denote

$$D = \{(x_0^{(1)}, x_1^{(1)}, \dots, x_n^{(1)}), \dots, (x_0^{(N)}, x_1^{(N)}, \dots, x_n^{(N)})\}.$$

Let the first N_1 cases to be fully, and for the remaining N_2 cases to be right censored in X_0 . In the example of unemployment duration, where X_0 denotes the duration of an unemployment spell for an individual and (X_1, \dots, X_n) a suitable vector of features/covariates, (right) censoring of X_0 may be due to the end of study or entrance in educational programs for unemployed. In the sequel, by “censored” it is meant right censored response.

The compound collective conditional likelihood (CCCL) of θ given the data

D is defined as

$$\begin{aligned}
CCCL_N(\theta) &= \prod_{j=1}^{N_1} p_\theta(x_0^{(j)} | x_{[n]}^{(j)}) \prod_{k=N_1+1}^N P_\theta(X_0 > x_0^{(k)} | x_{[n]}^{(k)}) \\
&= \prod_{j=1}^{N_1} p_\theta(x_0^{(j)} | x_{[n]}^{(j)}) \prod_{k=N_1+1}^N \sum_{x_0=x_0^{(k)}+1}^{r_0} p_\theta(x_0 | x_{[n]}^{(k)}) \\
&= \prod_{j=1}^{N_1} \frac{p_\theta(x_0^{(j)}) \prod_{i=1}^n p_\theta(x_i^{(j)} | x_0^{(j)})}{\sum_{x'_0} p_\theta(x'_0) \prod_{i=1}^n p_\theta(x_i^{(j)} | x'_0)} \\
&\quad \times \prod_{k=N_1+1}^N \sum_{x_0=x_0^{(k)}+1}^{r_0} \frac{p_\theta(x_0) \prod_{i=1}^n p_\theta(x_i^{(k)} | x_0)}{\sum_{x'_0} p_\theta(x'_0) \prod_{i=1}^n p_\theta(x_i^{(k)} | x'_0)} \\
&= \prod_{j=1}^{N_1} \frac{p_\theta(x_0^{(j)}) \prod_{i=1}^n p_\theta(x_i^{(j)} | x_0^{(j)})}{\sum_{x'_0} p_\theta(x'_0) \prod_{i=1}^n p_\theta(x_i^{(j)} | x'_0)} \\
&\quad \times \prod_{k=N_1+1}^N \left\{ \frac{P_\theta(X_0 = x_0^{(k)} + 1) \prod_{i=1}^n p_\theta(x_i^{(k)} | X_0 = x_0^{(k)} + 1)}{\sum_{x'_0} p_\theta(x'_0) \prod_{i=1}^n p_\theta(x_i^{(k)} | x'_0)} \right. \\
&\quad \left. + \dots + \frac{P_\theta(X_0 = r_0) \prod_{i=1}^n p_\theta(x_i^{(k)} | X_0 = r_0)}{\sum_{x'_0} p_\theta(x'_0) \prod_{i=1}^n p_\theta(x_i^{(k)} | x'_0)} \right\} \\
&= CCL_N^1(\theta) + \dots + CCL_N^M(\theta),
\end{aligned}$$

where $M = \prod_{j=N_1+1}^N (r_0 - x_0^{(j)})$. Thus, the CCCL is a sum of M collective conditional likelihoods.

The maximum compound collective conditional likelihood estimator (MCCCLE) is then

$$\hat{\theta}_N = \operatorname{argmax}_{\theta \in \Theta} \left\{ CCL_N^1(\theta) + \dots + CCL_N^M(\theta) \right\}.$$

MCCCLE has no closed form expression in general and need to be solved numerically. We will need below the following MCCLEs

$$\hat{\theta}_N^i = \operatorname{argmax}_{\theta \in \Theta} CCL_N^i(\theta) \quad \text{for } i = 1, \dots, M.$$

2 Strong consistency of MCCCLE

In this section, we give a proof for the strong consistency of MCCCLE. First, we need the following identifiability assumption as usual in maximum likelihood theory.

Assumption 2.1. (Identifiability Condition) *If $p_\theta(x_0 \mid x_{[n]}) = p_{\theta'}(x_0 \mid x_{[n]})$ for all x_0 and $x_{[n]}$ then $\theta = \theta'$.*

This condition requires that θ should be uniquely determined by the corresponding density $p_\theta(\cdot \mid \cdot)$. As shown earlier [4] for MCCLE, this does not always hold.

The collective conditional likelihood function for the naive Bayes network model with complete data is a concave down function in the parameters [2]. As noticed above, the compound collective conditional likelihood function with censored response is a sum of collective likelihood function.

Lemma 2.2. *The compound collective conditional likelihood function $CCCL_N(\theta)$ defined above is concave down in θ .*

Proof: First note that for any $\theta \in \Theta$, $p_\theta(x_0 \mid x_{[n]}) > 0$, and, therefore, so do the collective conditional likelihood functions composing $CCCL_N$. Furthermore, the sum of two convex functions with same support is convex and so do the sum of any number of convex functions with same domain, thereby yielding the result. □

Lemma 2.3. *If f and g are two convex functions on the same domain with their global minima at x_1 and x_2 respectively, then $f + g$ has its global minima at $tx_1 + (1 - t)x_2$ for some $t \in [0, 1]$.*

Proof: If $x_1 = x_2$ then the result holds for $t = 0$ or $t = 1$. Consider the case where $x_1 \neq x_2$. Then $\dot{f}(x_1) + \dot{g}(x_1) < 0$ and $\dot{f}(x_2) + \dot{g}(x_2) > 0$ or vice versa. Since $f + g$ is a convex function (sum of two convex functions with same domain), there must be a point x_3 such that $\dot{f}(x_3) + \dot{g}(x_3) = 0$, where $x_3 = tx_1 + (1 - t)x_2$ for some $t \in [0, 1]$. □

If $M = 2$ then, since $CCL_N^1(\theta)$ and $CCL_N^2(\theta)$ are concave down functions having their maxima at $\hat{\theta}_N^1$ and $\hat{\theta}_N^2$ respectively, $CCL_N^1(\theta) + CCL_N^2(\theta)$ which is also concave down has its maximum at $\hat{\theta}_N := t_D \hat{\theta}_N^1 + (1 - t_D) \hat{\theta}_N^2$ where t_D is a vector of the same length as θ^* , whose components are in $[0, 1]$, and which is dependent on the data D . Note that both $\hat{\theta}_N^1$ and $\hat{\theta}_N^2$ are consistent estimates for θ^* under Assumption 2.1 (since they are MCCLEs, see [4]).

$$P_{\theta^*} \left\{ \lim_{N_1 \rightarrow \infty} \hat{\theta}_N^1 = \theta^* \right\} = 1 \quad (1)$$

$$P_{\theta^*} \left\{ \lim_{N_1 \rightarrow \infty} \hat{\theta}_N^2 = \theta^* \right\} = 1 \quad (2)$$

In the sequel, we write $\theta = 0$ for $\theta \in \Theta^0$ to mean that all the components of θ are zeros and similarly for any inequality on θ . Now we can prove the strong consistency of MCCCLE.

Theorem 2.4. *Under Assumption 2.1, MCCCLE $\hat{\theta}_N$ is strongly consistent as follows,*

$$P_{\theta^*} \left\{ \lim_{N_1 \rightarrow \infty} \hat{\theta}_{M+N_1} = \theta^* \right\} = 1, \forall M. \quad (3)$$

Proof: We prove the result by induction. Let $M = 2$, then $\hat{\theta}_N = t_D \hat{\theta}_N^1 + (1 - t_D) \hat{\theta}_N^2 = \hat{\theta}_N^2 + t_D(\hat{\theta}_N^1 - \hat{\theta}_N^2)$. By (1) and (2) we have

$$P_{\theta^*} \left\{ \lim_{N_1 \rightarrow \infty} \hat{\theta}_N^1 = \theta^* \cap \lim_{N_1 \rightarrow \infty} \hat{\theta}_N^2 = \theta^* \right\} = 1. \quad (4)$$

$$(5)$$

Since $\hat{\theta}_N^1, \hat{\theta}_N^2$ and θ^* are finite and $0 \leq t_D \leq 1$ (therefore $0 \leq \limsup_{N \rightarrow \infty} t_D \leq 1$) we can write

$$P_{\theta^*} \left\{ \limsup_{N_1 \rightarrow \infty} t_D(\hat{\theta}_N^1 - \hat{\theta}_N^2) = 0 \right\} = 1 \quad (6)$$

$$P_{\theta^*} \left\{ \limsup_{N_1 \rightarrow \infty} \hat{\theta}_N^2 + t_D(\hat{\theta}_N^1 - \hat{\theta}_N^2) = \theta^* \right\} = 1 \quad (7)$$

$$P_{\theta^*} \left\{ \limsup_{N_1 \rightarrow \infty} \hat{\theta}_N = \theta^* \right\} = 1 \quad (8)$$

Similarly we can write

$$P_{\theta^*} \left\{ \liminf_{N_1 \rightarrow \infty} \hat{\theta}_N = \theta^* \right\} = 1. \quad (9)$$

Hence

$$P_{\theta^*} \left\{ \lim_{N_1 \rightarrow \infty} \hat{\theta}_N = \theta^* \right\} = 1 \quad (10)$$

Now assume that for $M > 2$, $\hat{\theta}_N := w_D^1 \hat{\theta}_N^1 + \dots + w_D^M \hat{\theta}_N^M$ where $w_D^1 + \dots + w_D^M = 1$, the maximizer of $CCCL_N(\theta)$, is a consistent estimator of θ^* . Then, for the case

of $M + 1$, assume for simplicity that the additional new censored observation is $x_0^{(N+1)} = r_0 - 2$. Then,

$$\begin{aligned} CCCL_{N+1}(\theta) &= \left\{ CCL_N^1(\theta) + \dots + CCL_N^M(\theta) \right\} \\ &\quad \times \left\{ \frac{P_\theta(X_0 = r_0 - 1) \prod_{i=1}^n p_\theta(x_i^{(N+1)} | X_0 = r_0 - 1)}{\sum_{x'_0} p_\theta(x'_0) \prod_{i=1}^n p_\theta(x_i^{(N+1)} | x'_0)} \right. \\ &\quad \left. + \frac{P_\theta(X_0 = r_0) \prod_{i=1}^n p_\theta(x_i^{(N+1)} | X_0 = r_0)}{\sum_{x'_0} p_\theta(x'_0) \prod_{i=1}^n p_\theta(x_i^{(N+1)} | x'_0)} \right\}. \end{aligned}$$

Let us rewrite this (with obvious new notation)

$$\begin{aligned} CCCL_{N+1}(\theta) &= \left\{ CCL_{N+1}^{a,1}(\theta) + \dots + CCL_{N+1}^{a,M}(\theta) \right\} \\ &\quad + \left\{ CCL_{N+1}^{b,1}(\theta) + \dots + CCL_{N+1}^{b,M}(\theta) \right\} \end{aligned}$$

Now denote $\hat{\theta}_{N+1}^a := w_D^{a,1} \hat{\theta}_{N+1}^{a,1} + \dots + w_D^{a,M} \hat{\theta}_{N+1}^{a,M}$ and $\hat{\theta}_{N+1}^b := w_D^{b,1} \hat{\theta}_{N+1}^{b,1} + \dots + w_D^{b,M} \hat{\theta}_{N+1}^{b,M}$ where $\sum_j w_D^{i,j} = 1$ for $i = a, b$, the maximizers of the first and second sums of CCLs respectively. By assumption they are consistent estimators of θ^* . Now similarly to the case $M = 2$, we can write $\hat{\theta}_{N+1} := u_D \hat{\theta}_{N+1}^a + (1 - u_D) \hat{\theta}_{N+1}^b$, where $0 \leq u_D \leq 1$, the maximizer of $CCCL_{N+1}(\theta)$, and show that it is strongly consistent for θ^* .

□

Corollary 2.5. $p_{\hat{\theta}_N}(x_0 | x_{[n]})$ is strongly consistent estimator of $p_{\theta^*}(x_0 | x_{[n]})$ for each $x_{[n]}$.

Proof: Immediate from the theorem since the densities $p_\theta(x_0 | x_{[n]})$ for all $x_{[n]}$ are rational functions of the parameter which have no poles in Θ^o .

□

ACKNOWLEDGEMENTS. The authors gratefully acknowledge the financial support of the Swedish Research Council (through the Swedish Initiative for Microdata Research in the Medical and Social Sciences (SIMSAM) and Ageing and Living Condition Program) and Swedish Research Council for Health, Working Life and Welfare (FORTE).

References

- [1] T. Neocleous and S. Portnoy, A Partially Linear Censored Quantile Regression Model for Unemployment Duration, *IRISS Working Paper Series*, **2008-07** (2008). <http://iriss.ceps.lu/documents/irisswp91.pdf>
- [2] T. Ross, H. Grunwald, P. Myllymaki and H. Tirri, On Discriminative Bayesian Network Classifiers and Logistic Regression, *Machine Learning*, **59**, (2005), 267 - 296.
- [3] G. Webb, J. Boughtton and Z. Wong, Not So Naive Bayes: Aggregating One-dependence Estimators, *Machine Learning*, **58**, (2005), 5 - 24.
- [4] P. Wijayatunga and S. Mase, Asymptotic Properties of Maximum Collective Conditional Likelihood Estimator for Naive Bayes Classifiers, Accepted in *International Journal of Statistics and Systems*, (2007). Related research report: <http://www.is.titech.ac.jp/research/research-report/B/B-432.pdf>.
- [5] P. Wijayatunga, S. Mase and M. Nakamura, Appraisal of Companies with Bayesian Networks, *International Journal of Business Intelligence and Data Mining*, **1**(3), (2006), 326 - 346.